

CLAIMS

What is claimed is:

5

1. A method of issuing a pair of ordered requests into a communication fabric, the communication fabric comprising a plurality of nodes interconnected by a plurality of point-to-point links, the method comprising the acts of:

10

issuing, from a source and into the communication fabric, a first request of the pair of ordered requests, the first request being directed to a first node of the plurality of nodes;

receiving the first request at the first node;

15

issuing a first response directed to the source, the first response acknowledging the receipt of the first request at the first node; and

20

issuing, from the source and into the communication fabric, a second request of the pair of ordered requests in response to the receipt of the first response.

2. The method as recited in claim 1, comprising the act of:

25

ordering the first request at the first node, wherein the first response acknowledges completion of the ordering.

3. The method as recited in claim 2, wherein the first request comprises a first address, and wherein the act of ordering the first request comprises the act of queuing the first request with respect to a request subsequently received by the first node, the subsequently received request comprising the first address.

5

4. The method as recited in claim 1, wherein the second request is directed to a second node of the plurality of nodes.

10

5. The method as recited in claim 1, wherein the first node comprises a first memory controller, wherein the second node comprises a second memory controller, and wherein the first request and second request are directed to the first memory controller and the second memory controller, respectively.

15

6. The method as recited in claim 1, wherein the source comprises a host bridge connected between the communication fabric and an input/output (I/O) subsystem, the I/O subsystem comprising a plurality of I/O nodes, and the method comprises the acts of:

20

generating a first I/O request by an I/O node of the plurality of I/O nodes, the first I/O request being directed to the host bridge;

generating a second I/O request by an I/O node of the plurality of I/O nodes, the second I/O request being directed to the host bridge;

25

receiving the first I/O request at the host bridge;

receiving the second I/O request at the host bridge;

translating the first I/O request into the first request; and

translating the second I/O request into the second request.

5

7. The method as recited in claim 6, comprising the acts of:

determining whether the second I/O request is ordered with respect to the first I/O
10 request; and,

if so,

stalling the act of issuing of the second request until the first response is
15 received by the host bridge.

8. The method as recited in claim 1, wherein the first request comprises a first
write request addressed to a first memory location in the communication fabric, and the
20 second request comprises a second write request addressed to a second memory location in
the communication fabric.

9. The method as recited in claim 8, wherein the first memory location is
25 accessible by the first node of the plurality of nodes, and the second memory location is
accessible by a second node of the plurality of nodes.

10. The method as recited in claim 1, wherein the first request comprises a write request addressed to a first memory location in the communication fabric, and the second request comprises a read request addressed to a second memory location in the communication fabric.

5

11. The method as recited in claim 10, wherein the first memory location is accessible by the first node of the plurality of nodes, and the second memory location is accessible by a second node of the plurality of nodes.

10

12. A method of issuing a plurality of requests from an input/output (I/O) system into a processing system, the processing system comprising a plurality of nodes interconnected by a plurality of point-to-point links, the processing system being connected to the I/O system via a host bridge, the method comprising the acts of:

15

issuing by the host bridge, into the processing system, a first request received from the I/O system, the first request being directed to a first node of the plurality of nodes;

20

determining by the host bridge, whether a second request received from the I/O system is ordered with respect to the first request; and,

if so,

25

stalling, by the host bridge, issuing of the second request;

receiving, at the host bridge, an acknowledgement that the first node has received the first request; and

issuing, by the host bridge, into the processing system, the second request in response to receipt of the acknowledgement.

5

13. The method as recited in claim 12, wherein if the host bridge determines the second request is not ordered with respect to the first request, the method comprises the act of issuing, by the host bridge, the second request into the processing system.

10

14. The method as recited in claim 12, comprising the acts of:

receiving, at the first node, the first request;

15

ordering, at the first node, the first request;

generating the acknowledgement by the first node, the acknowledgement indicating completion of the act of ordering the first request at the first node.

20

15. The method as recited in claim 14, wherein the first request comprises a first memory address, and wherein the act of ordering the first request at the first node comprises the act of queuing the first request with respect to other requests comprising the first memory address which are received at the first node subsequent to the first request, such that servicing of the first request occurs before servicing of the other requests.

25

16. The method as recited in claim 12, wherein the first node comprises a memory controller, and the first request comprises a first memory address controlled by the memory controller.

5

17. The method as recited in claim 12, wherein the first request is a write request.

10

18. The method as recited in claim 17, wherein the second request is a write request.

15

19. The method as recited in claim 12, wherein the second request is directed to a second node of the plurality of nodes.

20

20. The method as recited in claim 17, wherein the second request is a read request.

25

21. A computing system comprising:

a plurality of nodes;

a plurality of point-to-point links interconnecting the plurality of nodes to form a communication fabric;

a source to issue a plurality of requests into the communication fabric, wherein the source is configured to:

issue a first request of the plurality of requests into the communication fabric,
the first request being directed to a first node of the plurality of nodes;

5 to determine whether a second request of the plurality of requests is ordered
with respect to the first request; and,

if so,

10 to stall issuing the second request until the first request has been
received and ordered by the first node.

15 22. The system as recited in claim 21, wherein if the source determines that the
second request is not ordered with respect to the first request, then the source issues the
second request into the communication fabric.

20 23. The system as recited in claim 21, wherein the second request is directed to a
second node of the plurality of nodes.

25 24. The system as recited in claim 21, wherein the first request is a first write
request directed to a first memory location associated with the first node.

25 25. The system as recited in claim 24, wherein the second request is a second
write request directed to a second memory location associated with a second node of the
plurality of nodes.

26. The system as recited in claim 21, comprising an input/output (I/O) system coupled to the source, wherein the source is configured to receive a first I/O request and a second I/O request from the I/O system and to translate the first I/O request and the second I/O request into the first request and the second request, respectively.

27. The system as recited in claim 21, wherein the source comprises a host bridge.

28. The system as recited in claim 21, wherein the system comprises a memory, the memory being distributed into a plurality of memory portions, each of the memory portions being coupled to a respective node of the plurality of nodes.

29. The system as recited in claim 28, wherein the first node comprises:

an interface coupled to a first point-to-point link of the plurality of point-to-point links, the interface configured to receive the first request;

processing logic coupled to the interface to process the first request;

a cache coupled to the processing logic, the cache configured to store data and executable instructions;

a processor coupled to the cache to execute the instructions stored therein; and

a memory controller coupled to the processing logic to control access to the memory portion coupled to the first node, the memory controller comprising a request queue for ordering the first request.

5

30. The system as recited in claim 29, wherein the first request is directed to a first memory location associated with the memory portion coupled to the first node, and wherein the memory controller orders the first request with respect to other requests directed to the first memory location, the other requests being received subsequent to the first request.

10

31. The system as recited in claim 21, wherein the first node is configured to receive and order the first request and to generate a first response directed to the source, the first response acknowledging the receipt and ordering of the first request by the first node.

15

32. A computing system, comprising:

a coherent distributed communication fabric comprising at least a first target and a second target;

20

a non-coherent communication fabric; and

a bridge coupled between the coherent distributed communication fabric and the non-coherent communication fabric, wherein the bridge is configured to:

25

receive a pair of ordered requests from the non-coherent communication fabric, the pair of ordered requests comprising a first request directed to a first target in the coherent distributed communication fabric and a

second request directed to a second target in the coherent distributed communication fabric;

translate the first request and the second request into a first coherent request
and a second coherent request, respectively;

issue the first coherent request into the coherent distributed communication fabric; and

delay the second coherent request pending receipt of a first response from the first target; and

wherein the first target is configured to:

receive the first coherent request;

order the first coherent request with respect to other requests received by the first target; and

issue a first response directed to the bridge, the first response acknowledging receipt and ordering of the first coherent request.

33. The system as recited in claim 32, comprising a memory distributed into at least a first memory portion under control of the first target and a second memory portion under control of the second target.

34. The system as recited in claim 33, wherein the first target comprises:
a first memory controller coupled to the first memory portion; and
a queue to order the first request.

5

35. The system as recited in claim 32, wherein the non-coherent communication fabric comprises a plurality of input/output (I/O devices) interconnected by a plurality of point-to-point links.

10

36. A method of issuing a plurality of requests into a coherent communication fabric, the method comprising the acts of:

15

issuing, by a source, into the coherent communication fabric, a first request, the first request being directed to a first target in the coherent communication fabric;

determining whether a second request directed to a second target in the coherent communication fabric is ordered with respect to the first request; and,

20

if so,

stalling issuing of the second request into the coherent communication fabric pending receipt of an acknowledgement that the first request has been received and ordered at the first target.

25

37. The method as recited in claim 36, wherein if the second request is not ordered with respect to the first request, the method comprises the act of issuing the second request into the coherent communication fabric.

38. The method as recited in claim 36, comprising the acts of:

5 receiving the acknowledgement that the first request has been received and ordered at the first target; and

issuing the second request into the coherent communication fabric in response to the receipt of the acknowledgement.

10 39. The method as recited in claim 36, comprising the acts of:

receiving, at the first target, the first request; and

15 ordering, at the first target, the first request, wherein the first request is directed to a first address, and ordering comprises ordering the first request with respect to other requests directed to the first address which are subsequently received at the first target.

20 40. The method as recited in claim 39, comprising the act of;

25 generating the acknowledgement by the first target, the acknowledgement indicating completion of the act of ordering the first request at the first target; and

issuing, by the first target, the acknowledgement, the acknowledgement being directed to the source.

5

10

69